Enhanced Star Glyphs for Multiple-Source Data Analysis

Adrian Rusu[1], Confesor Santiago[2], Andrew Crowell[2], Eric Thomas[1]

[1]Department of Computer Science, Rowan University, Glassboro, NJ 08028, USA

[2]Federal Aviation Administration, Pomona, NJ 08240, USA

rusu@rowan.edu, {confesor.santiago, andrew.crowell}@faa.gov, thomas00@students.rowan.edu

Abstract

The analysis of large sums of data can be extremely difficult to perform if the data is not presented graphically. As a result, many graphing techniques have been developed, such as scatter plots, histograms. Generally, the main purpose of graphically displaying data is to do one of two things: First, to find the general average of where most of the data lies. Second, to find the outliers, the data points that are most distant from the others. Our visualization will attempt to find both by using a multitude of common graphing techniques to expand upon the traditional star glyph and create a new way of graphing data. These techniques include clustering, using color as identifiers, and 3D graphing capabilities to present more data that would not be possible of being shown in a two dimensional environment. We apply our techniques to compare several air traffic trajectory predictors currently being analyzed by the U.S. Federal Aviation Administration.

Keywords—**star glyph, air traffic, 3D, application**

1 Introduction

Representing data in different contexts can present a whole new meaning to the way it is interpreted. Many graphs have successfully been displayed in a two-dimensional environment. However, these graphs are limited to the amount of data that can be displayed at one time because of the constraints of physical viewing space a two-dimensional graph has. By adding in an extra axis, there are virtually infinitely many new ways to display data.

A star glyph is a multivariate graphing technique in which each variable represents a ray, or "spoke," each of which extends out via a connecting line from a common origin with equal angular distance between each spoke. The length of each line is proportional to the magnitude of the variable compared to the maximum value of all the variables [2]. Star glyphs are extremely effective at determining both where data begins to cluster and where the outliers are. However, the downfall occurs when the data set becomes increasingly large as the graph becomes overwhelmingly cluttered.

Many statistical graphing techniques are available to help overcome the clutter of large data sets. Scatter plots, parallel coordinates, and star plots are all have done a very good job at creating more modern techniques such as GeoTime [5], Jigsaw [9], NetLens [10] Substrate Designer [15], and Dust & Magnet [3].

This paper discusses a new way of graphing to visualize quantitative data in a three dimensional environment. We take the basic concept of a star glyph, a figure with n rays, where n is the dimensionality of the data set [17], and add a second attribute to the end of the spoke represented in the form of a sphere. By adding a third axis we overcome the cluttering that occurs when grouping large data sets. The result is a four dimensional multivariate plot, not limited to a two dimensional plane, useful for comparing similar grouped attributes across multiple data sets. As an example, we apply these properties of an enhanced star glyph to a visualization showing the accuracy of aircraft trajectory predictions across multiple systems used in the U.S. National Airspace System.

The rest of the paper is organized as follows. Our enhancements to the basic two dimensional star glyph are described in Section 2. In Section 3 we present an application of our techniques to air traffic data. Conclusion and future work are discussed in Section 4.

2 Enhanced Star Glyph

In this section, we describe the concept of the enhanced star glyph.

2.1 Basic Star Glyph

A very basic star glyph extending out from a common origin in equal angular distances is presented in Figure 1. A star glyph has two dimensions: a quantifiable dimension and a categorical dimension. The quantifiable dimension is a single value and the magnitude of this value is represented by the length of a spoke. The quantifiable value is grouped by categorical data. For instance, suppose the star glyph in Figure 1 is showing the average amount of food daily a dog in a particular kennel will eat. Assume

the data is grouped by the breed of the dog and *s4* represents all beagles, while *s2* represents all St. Bernards. We can easily see that, on average, a St. Bernard will consume much more food than a beagle and all other types of dogs, while a beagle will consume close to the least amount of food. This graph allows us to recognize an outlier either for exclusion, or for further analysis.
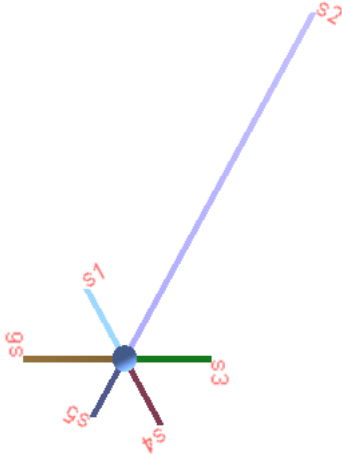

Figure 1: A Basic Star Glyph

2.2 An Additional Attribute

Our visualization has the feature of adding another attribute to compare more data. This attribute can be any arbitrary value represented in the form of sphere where the radius of the sphere is proportional to the attribute's value.

Consider the graph in Figure 2 representing the average number of miles per day an athlete runs grouped by the sport he or she plays. Let the size of the sphere represent the average weight of the athletes. Now, assume that the large green sphere to the bottom right represents American football players while the bottom purple sphere represents cross country runners. We can see that the cross country runners, on average, run much longer than football players, however, football players are generally much heavier.

The addition of the sphere allows for more data to be added to the graph without becoming overwhelming. Furthermore, the color of the spheres can put the data in a particular category. For example, in Figure 2 the two blue spheres on top and top-right represent lacrosse players. The sphere at the top represents lacrosse players in the fall, whereas the sphere on the top-right represents lacrosse players in the spring. From this graph we can see that lacrosse players run more in the spring since they are presumably better trained by that time.
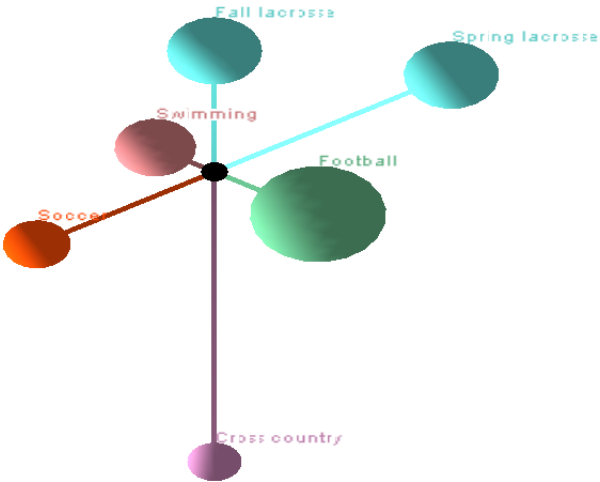

Figure 2: A Star Glyph with Spheres

Table 1 shows what a sample data set for the previous graph might look like.

Sport	Miles	Weight	Color
American football	1.5	250	Green
Cross country	8	130	Purple
American soccer	4	160	Red
Swimming	0.5	140	Orange
Fall lacrosse	3	150	Blue
Spring lacrosse	5	150	Blue

Table 1: Data Chart for the Previous Graph

2.3 Clustered Mean Graph

By adding the spheres to represent another measurement, we increase the clutter of a graph that has many points. This creates a cluster graph as shown in Figure 3, making it difficult to analyze the groups inside the cluster, but easily identifying and analyzing the outliers of the graph. In order to view these outliers we give an option of graphing a clustered mean star glyph.

Wherever data is closely similar, the visualization clusters that data into a "blob" of spheres. It removes the labels for clustered data and leaves the labels for the outliers.

The average data are always clustered in the center and the outliers extend out in a position that represents their absolute value from the average cluster.
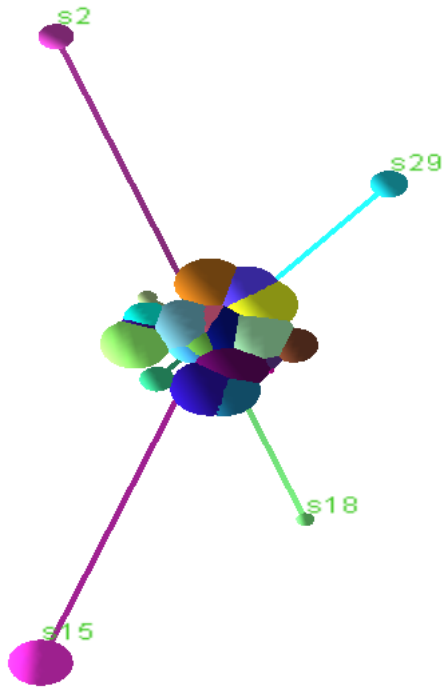
Figure 3: A Clustered Star Glyph

2.4 Comparison of Multiple Star Glyphs

Finally, we add the last dimension of our visualization. A single star glyph exists on a single two dimensional plane, therefore, since this is a three dimensional graph, we have the third dimension to display multiple star glyphs. This allows us to compare several results to each other.

For instance, consider that Figure 4 represents surveys taken by several different news agencies on which of six different companies produces the best ice cream. The distance of the spokes represents the average score of the ice cream producer (from 0 to 5, 5 being best), while the size of the sphere represents the amount of sales that producer had last year. Each separate glyph represents a separate survey. By displaying the three different surveys in same visualization, we can determine, among other things, which survey was more biased toward a particular producer.
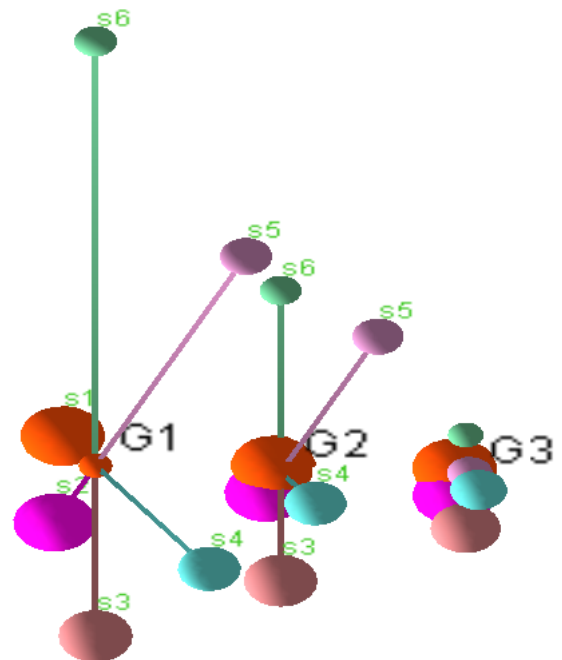

Figure 4: Multiple Star Glyphs

We use the following algorithm to determine whether or not data should be included in the cluster.

Algorithm Clustering

Input: An input file containing the data points of the nodes.
Output: A graph with nodes. Nodes in the graph are clustered if they are close to the average of the data points. Nodes that are not clustered are the outliers and their positions are the absolute values from the average of the nodes.

for each $node$ in $list\_of\_nodes$
   calculate the avg of all the nodes
for each $node$ in $list\_of\_nodes$
      if $node$ is close to avg **then**
      add $node$ to cluster

end Algorithm.

Clustering data is increasingly more important in analysis especially in a topic such as data mining. By abstracting the details of the average, an analyst can focus only on the outliers. In data mining, an outlier could be a group of people that has not been marketed to or fraudulent data in a company's budget.

2.5 User Interaction

Allowing the user to easily navigate the scene to focus on points of interest would give him invaluable insight into the information presented. One basic interaction is the ability to rotate the star glyphs in any direction, allowing the user to get the perspective he desires (see Figure 5), and to move about the visualization to find the exact camera position required for analysis of a focal point.
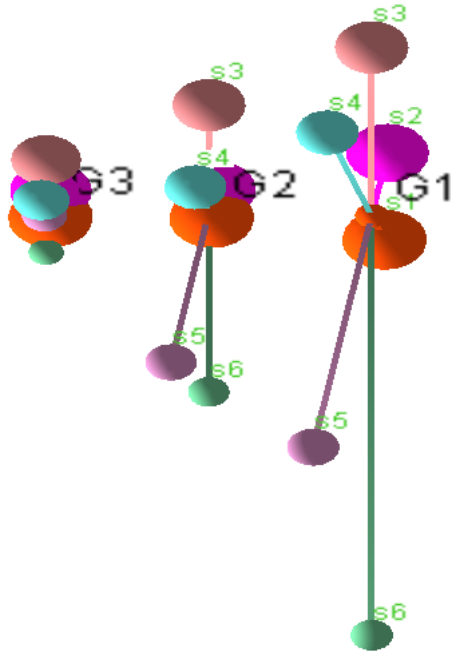
Figure 5: Multiple Star Glyphs at a Different Angle

Another interaction detail is allowing the user to manipulate the scale of the plot along each axis. This includes moving the planar star glyphs closer together for a more clustered appearance, or spreading them out more (see Figure 6). Grouping the data in this way visually displays where data begins to cluster across the multiple data sets. This helps analysts discover which of the outliers of the data set are more distant from the other outliers.
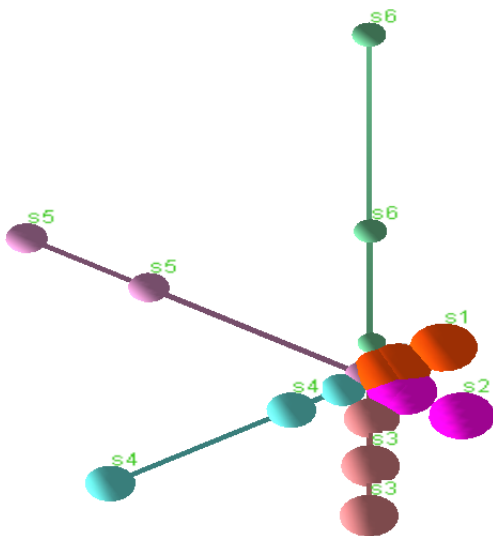

Figure 6: Grouped Data Sets

Another necessity is the ability to change the scales of both the spokes and the spheres. This can make a point of interest less cluttered or easier to distinguish.

In many visualizations, labeling of nodes can cause cluttering and are often indistinguishable from each other. Giving the user the ability to display and hide labels of his choosing is a necessary provision. Each spoke, sphere, and planar glyph has its own label and, often, have many more details that are too much to display in the visualization. For instance, each sphere may be labeled with what category it represents, but the user may be interested in the exact value the size of the sphere represents. By clicking on that particular sphere, detailed information about that category will be displayed, separate from the visualization. Using this feature, the user can view specific details for focal points to compare and analyze.

3 Application

We have implemented our visualization techniques in Java using JOGL, OpenGL bindings, to generate the 3D graphics. In this section we present an application of our visualization.

3.1 Background

The U.S. Federal Aviation Administration (FAA) constantly strives for a safer, more efficient National Airspace System. It is predicted that by the year 2025, the current air traffic in the National Airspace System will have tripled what it is today. In order to handle this massive increase in air traffic density, many projects are under development as part of the FAA's Next Generation Air Transportation System [4]. In this system, the air traffic controllers will use Decision Support Tools to assist them in managing air traffic. An important element of these tools is a Trajectory Predictor (TP).

A TP assists the air traffic controller, by generating trajectories, four dimensional (latitude, longitude, altitude, time) predictions of where the aircraft will fly some time into the future. The trajectories are then processed by a Conflict Probe to predict when aircraft will come into conflict (i.e. loss of minimum separation standard; aircraft become too close together). A TP is not always accurate in predicting the future flight paths based on an array of factors (lack of aircraft intent, weather, etc.) Moreover, the accuracy of trajectories generated by a TP determines its overall performance [7].

There are currently two major Air Traffic Control systems in which a TP is a defining element. The User Request Evaluation Tool (URET) is currently deployed in all air traffic control facilities across the United States. Extensive research has been performed on URET in order to validate its algorithms and evaluate its accuracy [1], [11].

This air traffic control system is sufficient for now, but will quickly become out of date as air traffic increases. The solution is a new system called En Route Automation Modernization (ERAM). Currently being developed by the Lockheed Martin Corporation, this Air Traffic Control system will be the replacement for URET and will be more powerful and more efficient. The accuracy of the TP in ERAM has recently been analyzed and results are present in [8]. Furthermore, the FAA has several in-house TPs used for regression analysis and evaluation of ERAM as it is being developed.

One main part of regression analysis is the comparison of a dataset to another reference dataset. The comparison of TPs using an enhanced star glyph will allow analysts to quickly visualize strengths and weaknesses in each. This allows analysts and the ERAM developers to investigate where ERAM needs improvement, where other TPs are accurate, and how the current version of ERAM compares to a previous version.

3.2 Trajectory Predictor Comparison

For our example, we compared six different TPs: two major ones (ERAM and URET) as well as four different FAA in-house TPs, those being Linear Predictor, Flight Plan, Hybrid, and Hybrid Merge [6].

The TPs have a selection of algorithms that make predictions of aircraft based on a number of factors that can include weather, air traffic, aircraft type, flight plans, etc.

Our visualization graphically displays four different variable attributes for each TP. We will break down each variable so the reader is fully aware of how we determine the strengths and weaknesses of a TP [12].

- *Look Ahead Times* are the time intervals in the future that a TP makes a prediction at. Generally, time in the U.S. National Airspace System is measured in seconds UTC that have past so far in a given day. So 12:00AM (midnight) is considered 0 seconds, whereas 12:00PM (noon) is considered 43,200 seconds. A look ahead time is how many seconds into the future a TP will predict the position of the aircraft. There are six look ahead times, used in these particular datasets: 0, 60, 300, 600, 900, and 1200 seconds. What this means is we have a total look ahead time of 1200 seconds (20 minutes) and we analyze the accuracy of that 1200 second look ahead time at each listed time.

- *Engine Types* are large subsets that an aircraft falls under. The groups are based on the aircraft's engine which determines many things concerning the trajectory, including the top of ascent, speed, and distance

an aircraft can fly. There are three aircraft engine types: Piston(P), Turbo Prop(T), and Jet(J) [14].

- *Number of Measurements* determine the amount of data that is included in a subset. This can give an idea of how many flights are in the subset, how often a trajectory prediction was made, and even how long the flights in that subset were in the air.

In Table 2 we see eight rows of data. The second column refers to what type of aircraft it is. Here, J stands for Jet, P for Piston, and not shown is T for Turbo Prop. The last column is a list of look ahead times for URET pertaining to each of those aircraft types. The first column represents how many measurements were taken for a certain aircraft type and a certain look ahead time. For example, the first row indicates that in the database there are 4,737 measurements for jet aircraft types at a 0 seconds look ahead time. Row eight shows that there are 102 measurements for piston aircraft types at 60 seconds look ahead time. There are fewer measurements because there are less piston aircraft than jet aircraft in operation.

URET		
Count	**Engine Type**	**Look Ahead Time**
4737	J	0
4640	J	60
4258	J	300
3779	J	600
3297	J	900
2819	J	1200
100	P	0
102	P	60

Table 2: Sample Data Set for URET

- *Average Horizontal Error* is a statistical measurement used for comparing TPs. Horizontal error is the horizontal distance in nautical miles from the predicted location of the aircraft to the actual position of the aircraft at that predicted time. We take the average of all these values grouped by a single look ahead time and single engine type to generate one value in nautical miles for each grouping.

Table 3 shows a selection of data for ERAM. The second column shows how close, on average, the TP's prediction was to the actual value. A lower number indicates a higher level of accuracy. The first row shows that ERAM has made 202 predictions for aircrafts of piston type 600 seconds into the future. When those predictions are averaged together, the final result is 10.33.

ERAM			
Count	**Avg. Horz. Error**	**Engine Type**	**Look Ahead Time**
202	10.3289975	P	600
182	14.0510698	P	900
162	17.3151512	P	1200
354	.888160169	T	0
347	1.14047839	T	60
311	2.23915788	T	300
264	3.47411061	T	600
221	4.47824027	T	900
177	5.58823446	T	1200

Table 3: Sample Data Set for ERAM

The rest of the data shows that ERAM is better at predicting turbo prop aircraft than piston type aircraft. Also, as one might suspect, as the look ahead time grows larger, the average prediction becomes less accurate.

3.3 The Visualization

Finally, we can take all the data gathered for the six TPs and graph it visually using the enhanced star glyph.

The visualization presented in Figure 7, shows a front view of the TPs predictions at particular look ahead times and for different types of aircraft. The spokes of the graphs represent the average horizontal error. Each spoke is a TP's prediction at a particular look ahead time. The length of the spoke is the average horizontal error of the prediction. Spokes that extend further from the center indicate that the TP did a poor job at making its prediction.

The color of the spokes tells what kind of aircraft type a spoke is associated with. For this example, the pink color is for jet planes, blue is for turbo prop, and green is for piston.

Finally, the size of the spheres that extend off the spokes relate to the number of measurements that have been taken. Larger spheres signify more measurements. The pink spokes at the bottom of the graph have much larger spheres since jet planes are much more common than turbo prop or piston aircraft.

In Figure 7, the spokes at the top of the graph correspond to piston aircraft (the blue color is associated with piston aircraft). This is the current focus of the graph (the graph can easily be manipulated to focus on a different type of aircraft). We find that ERAM (star glyph on the far right) still needs work in predicting the positions of piston aircraft. This can be seen as the length of the spokes at the top are much more distant from the center for ERAM than the other TPs. Our visualization is confirmed when we look at the compiled version of piston aircraft at a 1200 second look ahead time in the following data set.

In Table 4, we see a set of data for all the TPs for piston aircraft at a 1200 second look ahead time. ERAM's average horizontal error is 17.32 nautical miles away from the actual value, which is much worse than the other TPs. Being able to predict the position of aircraft as early as possible is crucial in preventing aircraft from coming in conflict.

TP Type	**Count**	**Avg. Horz. Error**
URET	74	5.64460946
Linear Predictor	172	12.9068221
Flight Plan	171	5.939322398
Hybrid	172	10.0111151
Hybrid Merge	171	6.7263614
ERAM	162	17.3151512

Table 4: Compiled Data Set

Analysts and developers can now use this information to analyze, in further detail, the trajectories and determine where in ERAM the algorithm for piston aircraft is causing inaccurate predictions [13].

ERAM, however, is very accurate in its predictions of Jet aircraft, as we can see by the pink spheres at the bottom of Figure 7. This extreme difference in its accuracy is most likely due to the developers focusing mainly on jet aircraft for now, since jets are much more prevalent in the U.S. National Airspace System.

4 Conclusion

In this paper we presented techniques to improve the basic star glyph. By adding a sphere at the end of each spoke on the star glyph, another attribute could be graphed. The color of the sphere also indicates a category that a piece of data belongs to. Our clustering effect reduces the clutter problem that the star glyph suffers from. Finally, by arranging a set of star glyphs along a certain axis, a multitude of data sets can be compared.
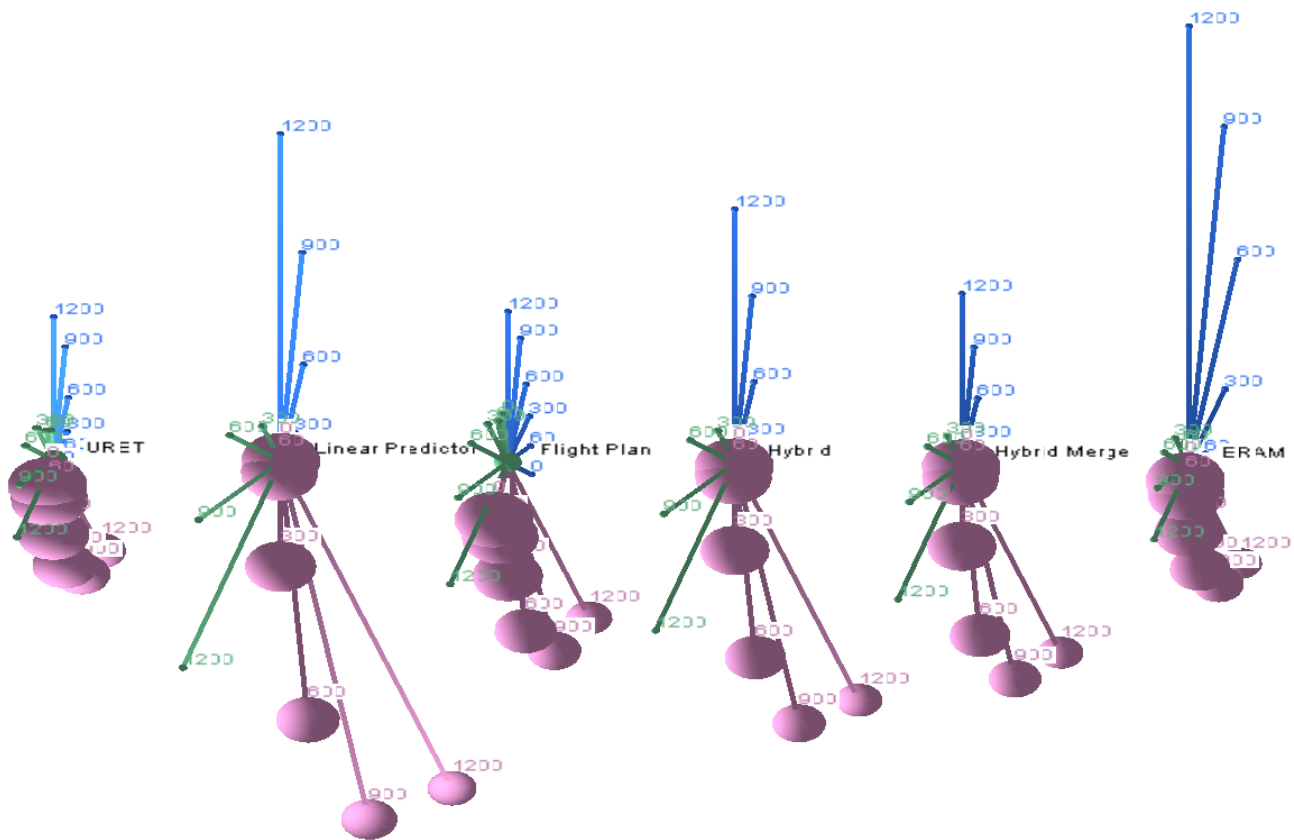
Figure 7: Multiple Enhanced Star Glyphs of Six Trajectory Predictors

We tested our techniques on six trajectory predictors currently being analyzed by the Federal Aviation Administration. We found that ERAM, the Air Traffic Control system set to be deployed within the next few years, currently performs well on horizontal prediction accuracy for jet and turboprop aircraft, but still needs work in predicting the horizontal locations of piston aircraft. This finding may indicate that ERAM needs to improve the performance model of this type of aircraft.

The data needs to be analyzed deeper in order to find the particular problem that is causing this. For this purpose a planar plot of just the points of interest would be useful. Our future work includes researching new ways to display the detailed information of the particular points the user is interested in investigating further.

Acknowledgments

References

[1] Mary Lee Cale, Stephen Kazunas, Mike M. Paglione, and Dr. Hollis F. Ryan. *User Request Evaluation Tool (URET) Algorithm Assessment Report.* DOT/FAA/CT-97/4. February 1997.

[2] John Chambers, William Cleveland, Beat Kleiner, and Paul Tukey. *Graphical Methods for Data Analysis.* Wadsworth. 1983.

[3] Ji Soo Yi, Rachel Melton, John Stasko and Julie A. Jacko. *Dust & Magnet: Multivariate Information Visualization Using a Magnet Metaphor.* Information Visualization Journal, pages 239-256, July 2005.

[4] FAA Fact Sheet - Next Gen `http://www.faa.gov/news/fact_sheets/news_story.cfm?newsId=8145`.

[5] Thomas Kapler and William Wright. *GeoTime Information Visualization.* Information Visualization Journal, pages 136-146, July 2005.

[6] Hollis F. Ryan and Mike M. Paglione. *State Vector Based Near Term Trajectory Prediction.* Proceedings American Institute of Aeronautics and Astronautics (AIAA) Guidance, Navigation, and Control Conference, Honolulu, Hawaii, August 2008.

[7] Hollis Ryan, Mike Paglione, and Steven Green. *Review of Trajectory Accuracy Methodology and Comparison of Error Measurement Metrics.* Proceedings American Institute of Aeronautics and Astronautics (AIAA) Guidance, Navigation, and Control Conference, Providence, Rhode Island, August 2004.

[8] Dr. Hollis F. Ryan, George Chandler, Confesor Santiago, Mike M. Paglione, and Shurong Liu. *Evaluation of En Route Host Computer Systems Trajectory Generation and Strategic Alert Processing: Analysis of ERAM Performance.* DOT/FAA/TC-TN08/10. November 2008.

[9] John Stasko, Carsten Görg and Zhicheng Liu. *Jigsaw: Supporting Investigative Analysis Through Interactive Visualization.* Information Visualization Journal, pages 118-132, April 2008.

[10] Hyunmo Kang, Catherine Plaisant, Bongshin Lee and Benjamin B. Bederson. *NetLens: Iterative Eploration of Content-Actor Network Data.* Information Visualization Journal, pages 18-31, March 2007.

[11] Mike M. Paglione, Dr. Hollis F. Ryan, Robert D. Oaks, J. Scott Summerill, and Mary Lee Cale. *Trajectory Prediction Accuracy Report User Request Evaluation Tool (URET) / Center-TRACON Automation System (CTAS).* DOT/FAA/CT-TN99/10. May 1999.

[12] Mike M. Paglione and Robert D. Oaks. *Implementation and Metrics for a Trajectory Prediction Validation Methodology.* Proceedings American Institute of Aeronautics and Astronautics (AIAA) Guidance, Navigation, and Control Conference, Hilton Head, South Carolina, August 2007.

[13] Confesor Santiago. *Using Graphical Software for Evaluating Aircraft Trajectory Predictions.* Proceedings Air Traffic Control Association Conference, 2005.

[14] US Sky Link `http://www.avchart.com/resource/charter-aircraft-types.asp`.

[15] Aleks Aris and Ben Shneiderman. *Designing Semantic Substrates for Visual Network Exploration.* Information Visualization Journal, pages 281-300, December 2007.

[16] Jonathan Roberts, Nadia Boukhelifa and Peter Rodgers. *Multiform Glyph Based Web Search Result Visualization.* Proceedings 6th International Conference on Information Visualisation, pages 549-554, IEEE, 2002.

[17] Matt Ward. *Xmdvtool: Integrating Multiple Methods for Visualizing Multivariate Data.* Proceedings IEEE Visualization, pages 326-333, 1994.